arXiv:1611.01421v1 [cs.CV] 4 Nov 2016

STDP-based spiking deep neural networks for object recognition

Saeed Reza Kheradpisheh[1,2,*], Mohammad Ganjtabesh[1], Simon J. Thorpe[2], and Timothée Masquelier[2]

[1] *Department of Computer Science, School of Mathematics, Statistics, and Computer Science, University of Tehran, Tehran, Iran*

[2] *CerCo UMR 5549, CNRS – Université Toulouse 3, France*

Abstract

Previous studies have shown that spike-timing-dependent plasticity (STDP) can be used in spiking neural networks (SNN) to extract visual features of low or intermediate complexity in an unsupervised manner. These studies, however, used relatively shallow architectures, and only one layer was trainable. Another line of research has demonstrated – using rate-based neural networks trained with back-propagation – that having many layers increases the recognition robustness, an approach known as deep learning. We thus designed a deep SNN, comprising several convolutional (trainable with STDP) and pooling layers. We used a temporal coding scheme where the most strongly activated neurons fire first, and less activated neurons fire later or not at all. The network was exposed to natural images. Thanks to STDP, neurons progressively learned features corresponding to prototypical patterns that were both salient and frequent. Only a few tens of examples per category were required and no label was needed. After learning, the complexity of the extracted features increased along the hierarchy, from edge detectors in the first layer to object prototypes in the last layer. Coding was very sparse, with only a few thousands spikes per image, and in some cases the object category could be reasonably well inferred from the activity of a single higher-order neuron. More generally, the activity of a few hundreds of such neurons contained robust category information, as demonstrated using a classifier on Caltech 101, ETH-80, and MNIST databases. We think that the combination of STDP with latency coding is key to understanding the way that the primate visual system learns, its remarkable processing speed and its low energy consumption. These mechanisms are also interesting for artificial vision systems, particularly for hardware solutions.

*Corresponding author.
Email addresses:
kheradpisheh@ut.ac.ir (SRK),
mgtabesh@ut.ac.ir (MG),
simon.thorpe@cnrs.fr (ST)
timothee.masquelier@cnrs.fr (TM).

1 Introduction

The primate's visual system solves the object recognition task through hierarchical processing along the ventral pathway of the visual cortex [1]. Through this hierarchy, the visual preference of neurons gradually increases from oriented bars in primary visual cortex (V1) to complex objects in inferotemporal cortex (IT), where neural activity provides a robust, invariant, and linearly-separable object representation [1, 2]. Despite the extensive feedback connections in the visual cortex, the first feed-forward wave of spikes in IT ($\sim 100-150$ ms post-stimulus presentation) appears to be sufficient

for crude object recognition [3, 4, 5].

During the last decades, various computational models have been proposed to mimic this hierarchical feed-forward processing [6, 7, 8, 9, 10]. Despite the limited successes of the early models [11, 12], recent advances in deep convolutional neural networks (DCNNs) led to high performing models [13, 14, 15]. Beyond the high precision, DCNNs can tolerate object variations as humans do [16, 17], use IT-like object representations [18, 19], and match the spatio-temporal dynamics of the ventral visual pathway [20].

Although the architecture of DCNNs is somehow inspired by the primate's visual system [21] (a hierarchy of computational layers with gradually increasing receptive fields), they totally neglect the actual neural processing and learning mechanisms in the cortex.

The computing units of DCNNs send floating-point values to each other which correspond to their activation level, while, biological neurons communicate to each other by sending electrical impulses (i.e., spikes). The amplitude and duration of all spikes are almost the same, so they are fully characterized by their emission time. Interestingly, mean spike rates are very low in the primate visual systems (perhaps only a few of hertz [22]). Hence, neurons appear to fire a spike only when they have to send an important message, and some information can be encoded in their spike times. Such spike-time coding leads to a fast and extremely energy-efficient neural computation in the brain (the whole human brain consumes only about 10-20 Watts of energy [23]).

The current top-performing DCNNs are trained with the supervised back-propagation algorithm which has no biological root. Although it works well in terms of accuracy, the convergence is rather slow because of the credit assignment problem [24]. Furthermore, given that DCNNs typically have millions of free parameters, millions of labeled examples are needed to avoid over-fitting. However, primates, specially humans, can learn from far fewer examples while most of the time no label is available. They may be able to do so thanks to Spike-timing-dependent plasticity (STDP), an unsupervised learning mechanism which occurs in mammalian visual cortex [25, 26, 27]. According to STDP, synapses through which a presynaptic spike arrived before (respectively after) a postsynaptic one are reinforced (respectively depressed).

To date, various spiking neural networks (SNNs) have been proposed to solve object recognition tasks. A group of these networks are actually the converted versions of traditional DCNNs [28, 29, 30]. The main idea is to replace each DCNN computing unit with a spiking neuron whose firing rate is correlated with the output of that unit. The aim of these networks is to reduce the energy consumption in DCNNs. However, the inevitable drawbacks of such spike-rate coding are the need for many spikes per image and the long processing time. Besides, the use of back-propagation learning algorithm and having both positive (excitatory) and negative (inhibitory) output synapses in a neuron are not biologically plausible. On the other hand, there are SNNs which are originally spiking networks and learn spike patterns. First group of these networks exploit learning methods such as auto-encoder [31, 32] and back-propagation[33] which are not biologically plausible. The second group consists of SNNs with bioinspired learning rules which have shallow architectures [34, 35, 36, 37, 38] or only one trainable layer [9, 39, 40].

In this paper we proposed a STDP-based spiking deep neural network (SDNN) with a spike-time neural coding. The network is comprised of a temporal-coding layer followed by a cascade of consecutive convolutional (feature extractor) and pooling layers. The first layer converts the input image into an asynchronous spike train, where the visual information is encoded in the temporal order of the spikes. Neurons in convolutional layers integrate input spikes, and emit a spike right after reaching their threshold. These layers are equipped with STDP to learn visual features. Pooling layers provide translation invariance and also compact the visual information [8]. Through the network, visual features get larger and more complex, where neurons in the last convolutional layer learn and detect object prototypes. At the end, a classifier detects the category of the input image based on the activity of neurons in the last pooling layer with global receptive fields.

We evaluated the proposed SDNN on Caltech face/motorbike and ETH-80 datasets with large-scale images of various objects taken form different viewpoints. The proposed SDNN reached the accuracies of 99.1% on face/motorbike task and 82.8% on ETH-80, which indicates its capability to recognize several natural objects even under severe variations. Based on our knowledge, there is no other spiking deep network which can recognize large-scale natural objects. We also examined the proposed SDNN on the MNIST dataset which is a benchmark for spiking neural networks, and interestingly, it reached 98.4% recognition accuracy. In addition to the high performance, the proposed SDNN is highly energy-efficient and works with a few number of spikes per image, which makes it suitable for neuromorphic hardware implementation.

2 Proposed Spiking Deep Neural Network

A sample architecture of the proposed SDNN with three convolutional and three pooling layers is shown in Figure 1. Note that the architectural properties (e.g., the number of layers and receptive field sizes) and learning parameters should be optimized for the desired recognition task.

The first layer of the network uses Difference of Gaussians (DoG) filters to detect contrasts in the input image. It encodes the strength of these contrasts in the latencies of its output spikes (the higher the contrast, the shorter the latency). Neurons in convolutional layers detect more complex features by integrating input spikes from the previous layer which detects simpler visual features. Convolutional neurons emit a spike as soon as they detect their preferred visual feature which depends on their input synaptic weights. Through the learning, neurons that fire earlier perform the STDP and prevent the others from firing via a winner-take-all mechanism. In this way, more salient and frequent features tend to be learned by the network. Pooling layers provide translation invariance using maximum operation, and also help the network to compress the flow of visual data. Neurons in pooling layers propagate the first spike received from neighboring neurons in the previous layer which are selective to the same feature. Convolutional and pooling layers are arranged in a consecutive order. Receptive fields gradually increase through the network and neurons in higher layers become selective to complex objects or object parts.

It should be noted that the internal potentials of all neurons are reset to zero before processing the next image. Also, learning only happens in convolutional layers and it is done layer by layer. Since the calculations of each neuron is independent of other adjacent neurons, to speed-up the computations, each of the convolution, pooling, and STDP operations are performed in parallel on GPU.

2.1 DoG and temporal coding

The important role of the first stage in SNNs is to encode the input signal into discrete spike events in the temporal domain. This temporal coding determines the content and the amount of information carried by each spike, which deeply affects the neural computations in the network. Hence, using efficient coding scheme in SNNs can lead to fast and accurate responses. Various temporal coding schemes can be used in visual processing (see ref. [41]). Among them, rank-order coding is shown to be efficient for rapid processing (even possibly in retinal ganglion cells) [42, 43].

Cells in the first layer of the network apply a DoG filter over their receptive fields to detect positive or negative contrasts in the input image. DoG well approximates the center-surround properties of the ganglion cells of the retina. When presented with an image, these DoG cells detect the contrasts and emit a spike; the more strongly a cell is activated (higher contrast), the earlier it fires. In other word, the order of the spikes depends on the order of the contrasts. This rank-order coding is shown to be efficient for obtaining V1 like edge detectors [44] as well as complex visual features [9, 40] in higher cortical areas.

DoG cells are retinotopically organized in two ON-center and OFF-center maps which are respectively sensitive to positive and negative contrasts. A DoG cell is allowed to fire if its activation is

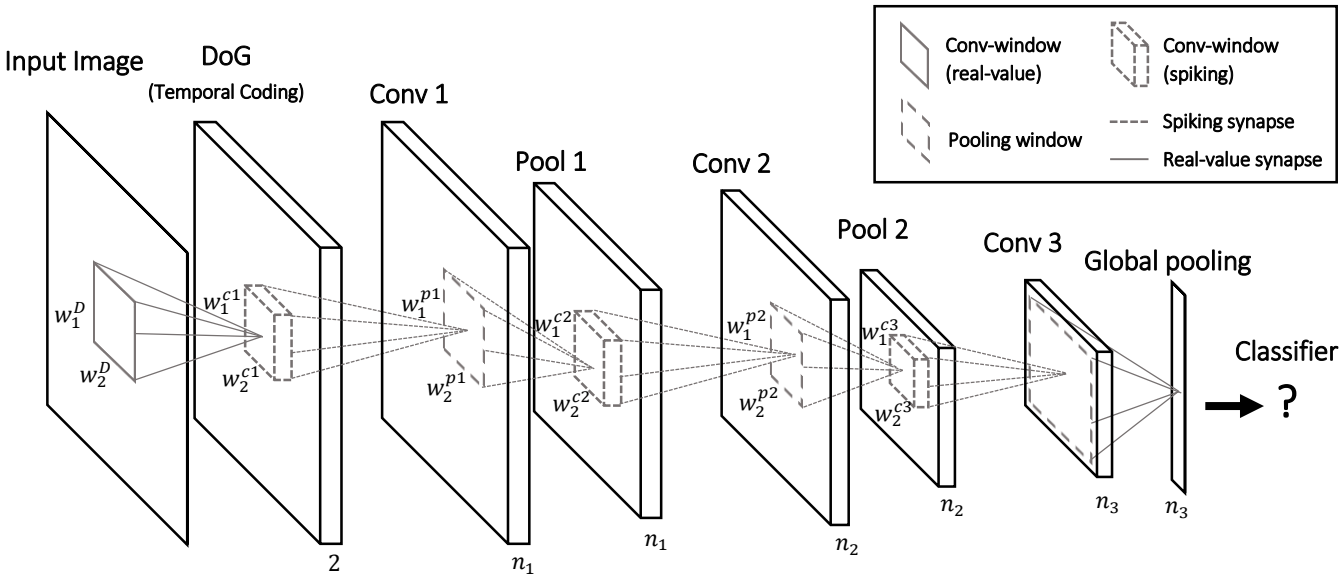

Figure 1: A sample architecture of the proposed SDNN with three convolutional and three pooling layers. The first layer applies ON- and OFF-center DoG filters of size $w_1^D \times w_2^D$ on the input image and encode the image contrasts in the timing of the output spikes. The ith convolutional layer, Conv i, learns combinations of features extracted in the previous layer. The ith pooling layer, Pool i, provides translation invariance for features extracted in the previous layer and compress the visual information using a local maximum operation. Finally the classifier detects the object category based on the feature values computed by the global pooling layer. The window size of the ith convolutional and pooling layers are indicated by $w_{1,2}^{ci}$ and $w_{1,2}^{pi}$, respectively. The number of the neuronal maps of the ith convolutional and pooling layer are also indicated by n_i below each layer.

above a certain threshold. Note that this scheme grantees that at most one of the two cells (positive or negative) can fire in each location. As mentioned above, the firing time of a DoG cell is inversely proportional to its activation value. For efficient GPU-based parallel computing, the input spikes are grouped into equal-size sequential packets. At each time step, spikes of one packet are propagated simultaneously. In this way, a packet of spikes with near ranks (carrying similar visual information) are propagated in parallel, while, the next spike packet will be processed in the next time step.

2.2 Convolutional layers

A convolutional layer contains several neuronal maps. Each neuron is selective to a visual feature determined by its input synaptic weights. Neurons in a specific map detect the same visual feature but at different locations. To this end, synaptic weights of neurons belonging to the same map should always be the same (i.e., weight sharing). Within a map, neurons are retinotopically arranged. Each neuron receives input spikes from the neurons located in a determined window in all neuronal maps of the previous layer. Hence, a visual feature in a convolutional layer is a combination of several simpler feature extracted in the previous layer. Note that the input windows of two adjacent neurons are highly overlapped. Hence, the network can detect the appearance of the visual features in any location.

Neurons in all convolutional layers are non-leaky integrate-and-fire neurons, which gather input spikes from presynaptic neurons and emit a spike when their internal potentials reach a prespecified threshold. Each presynaptic spike increases the neuron's potential by its synaptic weight. At each time step, the internal potential of the ith neuron is updated as follows:

$$V_i(t) = V_i(t-1) + \sum_j W_{j,i} S_j(t-1),$$

where $V_i(t)$ is the internal potential of the ith convolutional neuron at time step t, $W_{j,i}$ is the synaptic weight between the jth presynaptic neuron and the ith convolutional neuron, and S_j is the spike train of the jth presynaptic neuron ($S_j(t-1) = 1$ if the neuron has fired at time $t-1$, and $S_j(t-1) = 0$

otherwise). If V_i exceeds its threshold, V_{thr}, then the neuron emits a spike and V_i is reset:

$$V_i(t) = 0 \quad and \quad S_i(t) = 1, \quad if \quad V_i(t) \geq V_{thr}.$$

Also, there is a lateral inhibition mechanism in all convolutional layers. When a neuron fires, in an specific location, it inhibits other neurons in that location belonging to other neuronal maps and does not allow them to fire in the following time steps. In addition, neurons are not allowed to fire more than once. These together provides an sparse but highly informative coding, because, there can be at most one spike at each location which indicates the existence of a particular visual feature in that location.

2.3 Local pooling layers

Pooling layers help the network to gain invariance by doing a nonlinear max pooling operation over a set of neighboring neurons with the same preferred feature. Some evidence suggests that such a max operation occurs in complex cells in visual cortex [8]. Thanks to the rank-order coding used in the proposed network, the maximum operation of pooling layers simply consists of propagating the first spike emitted by the afferents [45].

A neuron in a neuronal map of a pooling layer performs the maximum operation over a window in the corresponding neuronal map of the previous layer. Pooling neurons are integrate-and-fire neurons whose input synaptic weights and threshold are all set to one. Hence, the first input spike activates them and leads to an output spike. Regarding to the rank-order coding, each pooling neuron is allowed to fire at most once. It should be noted that no learning occurs in pooling layers.

Another important role of pooling layers is to compress the visual information. Regarding to the maximum operation performed in pooling layers, adjacent neurons with overlapped inputs would carry redundant information (each spike is sent to many neighboring pooling neurons). Hence, in the proposed network, the overlap between the input windows of two adjacent pooling neurons (belonging to the same map) is set to be very small. It helps to compress the visual information by eliminating the redundancies, and also, to reduce the size of subsequent layers.

2.4 STDP-based learning

As mentioned above, learning occurs only in convolutional layers which should learn to detect visual features by combining simpler features extracted in the previous layer. The learning is done layer by layer, i.e., the learning in a convolutional layer starts when the learning in the previous convolutional layer is finalized. When a new image is presented, neurons of the convolutional layer compete with each other and those which fire earlier trigger STDP and learn the input pattern.

A simplified version of STDP [9] is used:

$$\begin{cases} \Delta w_{ij} = a^+ w_{ij}(1 - w_{ij}), & if \quad t_j - t_i \leq 0, \\ \Delta w_{ij} = a^- w_{ij}(1 - w_{ij}), & if \quad t_j - t_i > 0, \end{cases}$$

where i and j respectively refer to the index of post- and presynaptic neurons, t_i and t_j are the corresponding spike times, Δw_{ij} is the synaptic weight modification, and a^+ and a^- are two parameters specifying the learning rate. Note that the exact time difference between two spikes does not affect the weight change, but only its sign is considered. Also, it is assumed that if a presynaptic neuron does not fire before the postsynaptic one, it will fire later. These simplifications are equivalent to assuming that the intensity-latency conversion of DoG cells compresses the whole spike wave in a relatively short time interval (say, $20-30$ ms), so that all presynaptic spikes necessarily fall close to the postsynaptic spike time, and the time lags are negligible. The multiplicative term $w_{ij}(1-w_{ij})$ ensures the weights remain in the range [0,1] and thus maintains all synapses in an excitatory mode in adding to implementing soft-bound effect.

During the learning of a convolutional layer, neurons in the same map, detecting the same feature in different locations, integrate input spikes and compete with each other to do the STDP. The first neuron which reaches the threshold and fires, if any, is the winner (global intra-map competition). The winner triggers the STDP and updates its synaptic

weights. As mentioned before, neurons in different locations of the same map have the same input synaptic weights (i.e., weight sharing) to be selective to the same feature. Hence, the winner neuron prevents other neurons in its own map to do STDP and duplicates its updated synaptic weights into them.

Also, there is a local inter-map competition for STDP. When a neuron is allowed to do the STDP, it prevents the neurons in other maps within a small neighborhood around its location from doing STDP. This competition is crucial to encourage neurons of different maps to learn different features.

Because of the discretized time variable in the proposed model, it is probable that some competitor neurons fire at the same time step. One possible scenario is to pick one randomly and allow it to do STDP. But a better alternative is to pick the one which has the highest potential indicating higher similarity between its learned feature and input pattern.

Synaptic weights of convolutional neurons initiate with random values drown from a normal distribution with the mean of 0.8 and STD of 0.05. As the learning of a specific layer progresses, its neurons gradually converge to different visual features which are frequent in the input images. Also, the complexity of the visual features gradually increases through the layers in such a way that neurons in the highest layer converge to object prototypes.

2.5 Global pooling and classification

The global pooling layer is only used in the classification phase. Neurons of the last layer perform a global max pooling over their corresponding neuronal maps in the last convolutional layer. Such a pooling operation provides a global translation invariance for prototypical features extracted in the last convolutional layer. Hence, there is only one output value for each feature, which indicates the presence of that feature in the input image. The output of the global pooling layer over the training images is used to train a classifier, say an SVM. In the testing phase, the test object image is processed by the network and the output of the global pooling layer is fed to the classifier to determine its category.

To compute the output of the global pooling layer, first, the threshold of neurons in the last convolutional layer were set to be infinite, and then, their final potentials (after propagating the whole spike train generated by the input image) were measured. These final potentials can be seen as the number of early spikes in common between the current input and the stored prototypes in the last convolutional layer. Finally, the global pooling neurons compute the maximum potential at their corresponding neuronal maps, as their output value.

3 Results

3.1 Caltech face/motorbike dataset

We evaluated our SDNN on the face and motorbike categories of the Caltech 101 dataset available at http://www.vision.caltech.edu (see Figure 3 for sample images). The training set contains 200 randomly selected images per category, and remaining images constitute the test set. The test images are not seen during the learning phase but used afterward to evaluate the performance on novel images. This standard cross-validation procedure allows measuring the system's ability to generalize, as opposed to learning the specific training examples. All images were converted to grayscale values and rescaled to be 160 pixels in height (preserving the aspect ratio).

Here, we used a network similar to Figure 1, with three convolutional layers each of which followed by a pooling layer. For the first layer, only ON-center DoG filters of size 7×7 and standard deviations of 1 and 2 pixels are used. The first, second and third convolutional layers consists of 4, 20, and 10 neuronal maps with conv-window sizes of 5×5, $16 \times 16 \times 4$, and $5 \times 5 \times 20$ and firing thresholds of 10, 60, and 2, respectively. The pooling window sizes of the first and second pooling layers are 7×7 and 2×2 with the strides of 6 and 2, correspondingly. The third pooling layer performs a global max pooling operation. The learning rates of all convolutional

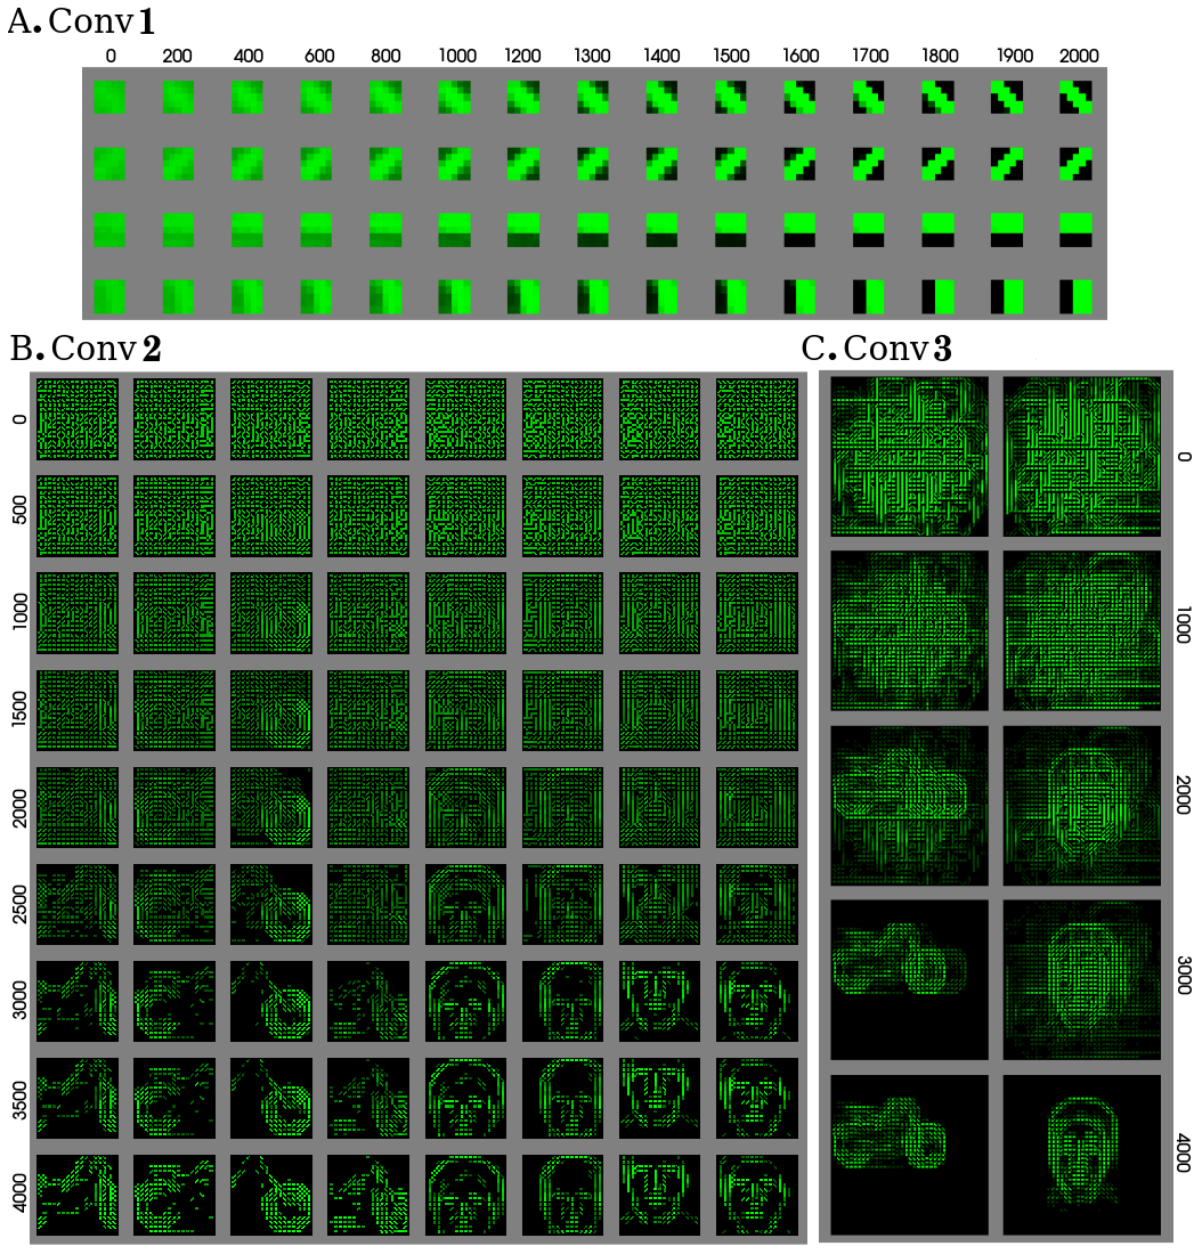

Figure 2: The synaptic changes of some neuronal maps in different layers through the learning with the Caltech face/motorbike dataset. A) The first convolutional layer becomes selective to oriented edges. B) The second convolutional layer converges to object parts. C) The third convolutional layer learns the object prototype and respond to whole objects.

layers are set to $a^+ = 0.004$ and $a^- = 0.003$. In addition, each image is processed for 30 time steps.

Figure 2 shows the preferred visual features of some neuronal maps in the first, second and third convolutional layers through the learning process. To visualize the visual feature learned by a neuron, a backward reconstruction technique is used. Indeed, the visual features in the current layer can be reconstructed as the weighted combinations of the visual features in the previous layer. This backward process continues until the first layer, whose preferred visual features are computed by DoG functions. As shown in Figure 2A, interestingly, each of the four neuronal maps of the first convolutional layer converges to one of the four orientations: $\pi/4$, $\pi/2$, $3\pi/4$, and π. This shows how efficiently the association of the proposed temporal coding in DoG cells and unsupervised learning method (the STDP and learning competition) led to highly diverse edge detectors which can represent the input image with

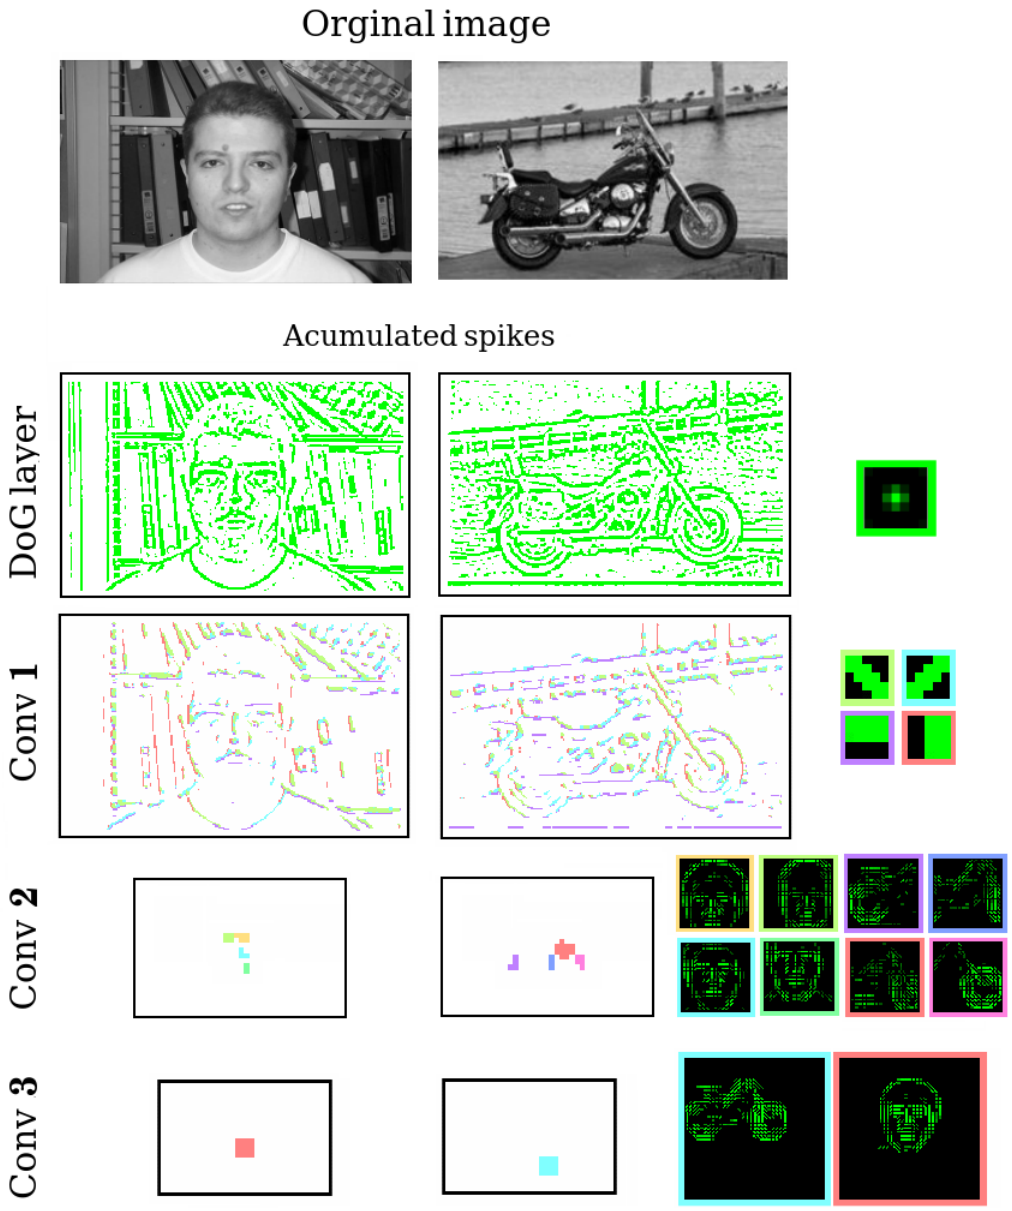

Figure 3: **The spiking activity of the convolutional layers with the face and motorbike images.** The preferred features of neuronal maps in each convolutional layer are shown on the right. Each feature is coded by a specific color border. The spiking activity of the convolutional layers, accumulated over all the time steps, is shown in the corresponding panels. Each point in a panel indicates that a neuron in that location has fired at a time step, and the color of the point indicates the preferred feature of the activated neuron.

edges in different orientations. These edge detectors are similar to the simple cells in primary visual cortex (i.e., V1 area) [44]

Figure 2B shows the learning progress for the neuronal maps of the second convolutional layer. As mentioned, the first convolutional layer detects edges with different orientations all over the image, and due to the used temporal coding, neurons corresponding to edges with higher contrasts (i.e., salient edges) will fire earlier. On the other hand, STDP naturally tends to learn those combination of edges that are consistently repeating in the training images (i.e., common features between the target objects). Besides, the learning competition tends to prevent neuronal the maps from learning similar visual features. Consequently, neurons in the second convolutional layer learn the most salient, common, and diverse visual features of the target objects, and do not learn the backgrounds that drastically change between images. As seen in Figure 2B, each of the maps gradually learns a different visual feature (combination of oriented

edges) representing a face or motorbike feature.

The learning progress for two neuronal maps of the third convolutional layer are shown in Figure 2C. As seen, one of them gradually becomes selective to a complete motorbike prototype as a combination of motorbike features such as back wheel, middle body, handle, and front wheel detected in the second layer. Also, the other map learns a whole face prototype as a combination of facial features. Indeed, the third convolutional layer learns the whole object prototypes using intermediate complexity features detected in the previous layer. Neurons in the second layer compete with each other and send spikes toward the third layer as they detect their preferred visual features. Since, different combinations of these features are detected for each object category, neuronal maps of the third layer will learn different prototypes of different categories. Therefore, the STDP and the learning competition mechanism direct neuronal maps of the third convolutional layer to learn highly category specific prototypes.

Figure 3 shows the accumulated spiking activity of the DoG and the following three convolutional layers over all time steps, for two face and motorbike sample images. For each layer, the preferred features of some neuronal maps with color coded borders are demonstrated on top, and their corresponding spiking activity are shown in panels below them. Each colored point inside a panel indicates the neuronal map of the neuron which has fired in that location at a time step. As seen, neurons of the DoG layer detect image contrasts, and edge detectors in the first convolutional layer detect the orientation of edges. Neurons in the second convolutional layer, which are selective to intermediate complexity features, detect their preferred visual feature by combining input spikes from edge detector cells in the first layer. Finally, the coincidence of these features activates neurons in the third convolutional layer which are selective to object prototypes. As seen, when a face (motorbike) image is presented, neurons in the face (motorbike) maps fire. To better illustrate the learning progress of all the layers as well as their spiking activity in the temporal domain, we prepared a short video (see Video S1 in Supplementary Information).

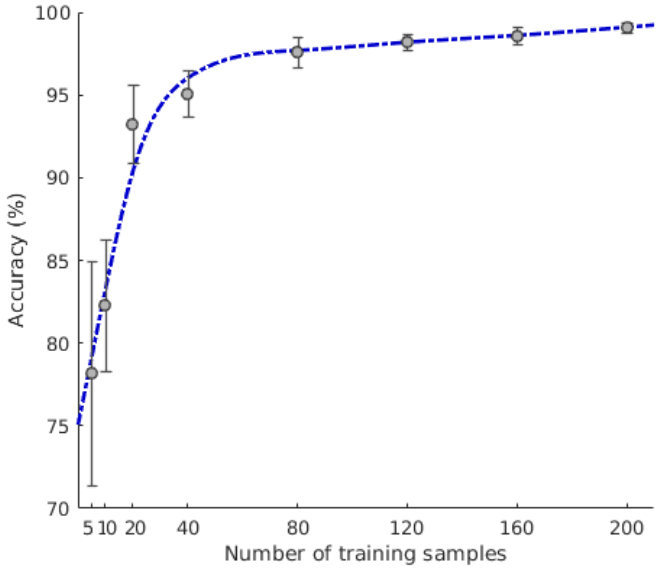

Figure 4: Recognition accuracies (mean ± std) of the proposed SDNN for different number of training images per category.

As mentioned in the previous section, the output of the global pooling layer is used by a SVM classifier to specify the object category of the input images. We trained the proposed SDNN on training images and evaluated it over the test images, where the model reached the categorization accuracy of $99.1 \pm 0.2\%$. It shows how the object prototypes, learned in the highest layer, can well represent the object categories. Furthermore, we also calculated the single neuron accuracy. In more details, we separately computed the recognition accuracy of each neuron in the global pooling layer. Surprisingly, some single neurons reached an accuracy of 93%, and the mean accuracy was 89.8%. Hence, it can be said that single neurons in the highest layer are highly class specific, and different neurons carry complementary information which altogether provide robust object representations.

In a further experiment, we changed the number of training samples (from 5 to 200 images per category) and calculated the recognition accuracy of the proposed SDNN. As seen in Figure 4, with 5 images per category it reached the accuracy of 78.2%, and only 40 images from each category are sufficient to reach 95.1% recognition accuracy. Although having more training samples leads to higher accuracies, the proposed SDNN can extract diagnostic

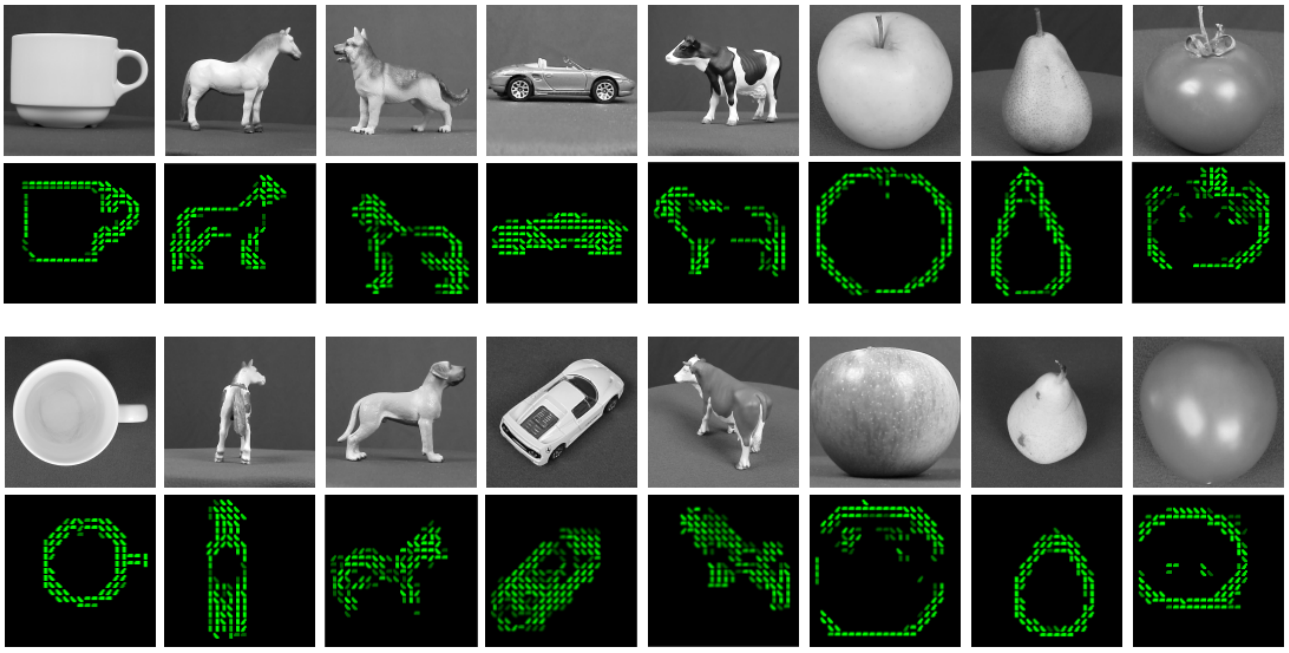

Figure 5: Some sample images of different object categories of ETH-80 in different viewpoints. For each image, the preferred feature of an activated neuron in the third convolutional layer is shown in below.

features and reach reasonable accuracies even using a few tens of training images. Due to the unsupervised nature of STDP, the proposed SDNN does not suffer much from the overfitting challenge caused by small training set size in supervised learning algorithms such as back-propagation.

Table 1: Recognition accuracies of the proposed SDNN and some other methods over the ETH-80 dataset.

Method	Accuracy (%)
HMAX [40]	69.0
Logistic regression [46]	77.0
Convolutional SNN [40]	81.1
Imagenet pre-trained DCNN [40]	79.5
Proposed SDNN	**82.8**

3.2 ETH-80 dataset

The ETH-80 dataset contains eight different object categories: apple, car, toy cow, cup, toy dog, toy horse, pear, and tomato (10 instances per category). Each object is photographed from 41 viewpoints with different view angles and different tilts. Some examples of objects in this dataset are shown in Figure 5. ETH-80 is a good benchmark to show how the proposed SDNN can handle multi-object categorization tasks with high inter-instance variability, and how it can tolerate huge view-point variations.

Five randomly chosen instances of each object category are selected for the training set used in the learning phase. The remaining instances constitute the testing set, and are not seen during the learning phase. All the object images were converted to grayscale values. To evaluate the proposed SDNN on ETH-80, we used a network architecturally similar to the one used for Caltech face/motorbike dataset. The other parameters are also similar, except for the number of neuronal maps in the second and third convolutional and pooling layers. Here we used 400 neuronal maps in each of these layers.

Similar to the caltech dataset, neuronal maps of the first convolutional layer converged to the four oriented edges. Neurons in the second and third convolutional layers also became selective to intermediate features and object prototypes, respectively. Figure 5 shows sample images from the ETH-80 dataset and the preferred features of some neuronal maps in the third convolutional layer which are activated for those images. As seen, neurons in the highest layer respond to different views of different objects, and altogether, provide an in-

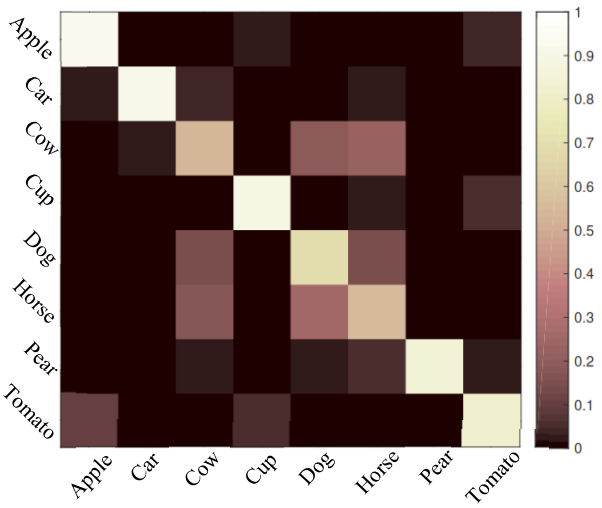

Figure 6: The confusion matrix of the proposed SDNN over the ETH-80 dataset.

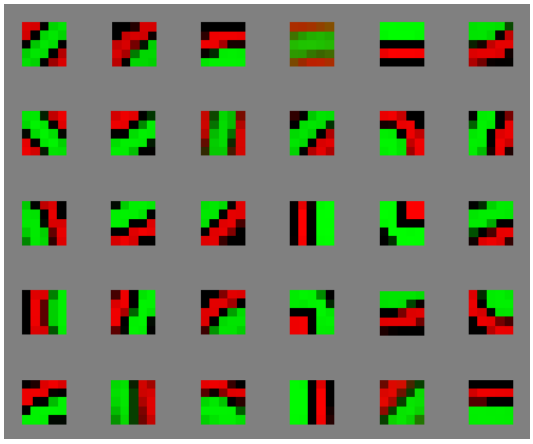

Figure 7: The Gabor-like features learned by the neuronal maps of the first convolutional layer from the MNIST images. The red and green colors receptively indicate the strength of input synapses from ON- and OFF-center DoG cells.

variant object representation. Thus, the network learns 2D and view-dependent prototypes of each object category to achieve 3D representations.

As mentioned before, we evaluated the proposed SDNN over the test instances of each object category which are not shown to the network during the training. The recognition accuracy of the proposed SDNN along with some recently reported results on ETH-80 dataset are presented in Table 1. Note that, the other methods were also trained and tested with the same protocol, five distinct instances per object category being used for each of the training and testing phases. The results indicate that the proposed SDNN outperforms the other methods with a recognition accuracy of 82.8%. This demonstrates the ability of the proposed SDNN to achieve an invariant object representation for multiple object categories.

In a subsequent analysis, we computed the confusion matrix, to see which categories are mostly confused with each other. Figure 6 illustrates the confusion matrix of the proposed SDNN over the ETH-80 dataset. As seen, most of the errors are due to the miscategorization of dogs, horses, and cows. We checked whether these errors belong to the some specific viewpoints or not. We found out that the errors are uniformly distributed between different viewpoints. Hence, it can be concluded that these categorization errors are due to the overall shape similarity between these object categories.

The other important aspect of the proposed SDNN is the computational efficiency of the network. For each ETH-80 image, on average, about 9100 spikes are emitted in all the layers, i.e., about 0.02 spike per neuron per image. This points to the fact that the proposed SDNN can recognize objects with high precision but low computational cost. This efficiency is caused by the association of the proposed temporal coding and STDP learning rule which led to a sparse but informative visual coding.

3.3 MNIST dataset

MNIST [47] is a benchmark dataset for SNNs which has been widely used [48, 37, 36, 49, 38, 50]. We also evaluated our SDNN on the MNIST dataset which contains 60,000 training and 10,000 test handwritten single-digit images. Each image is of size 28×28 pixels and contains one of the digits 0–9. For the first layer, ON- and OFF-center DoG filters with standard deviations of 1 and 2 pixels are used. The first and second convolutional layers respectively consist of 30 and 100 neuronal maps with 5×5 convolution-window and firing thresholds of 15 and 10. The pooling-window of the first pooling layer was of size 2×2 with the stride of 2. The second pooling layer performs a global max operation. Note that the learning rates of all convolutional layers were set to $a^+ = 0.004$ and $a^- = 0.003$.

Figure 7 shows the preferred features of some

Table 2: Recognition accuracies of the proposed SDNN and some other SNNs over the MNIST dataset.

Architecture	Neural coding	Learning-type	Learning-rule	Accuracy (%)
Dendritic neurons [48]	Rate-based	Supervised	Morphology learning	90.3
Convolutional SNN [37]	Spike-based	Supervised	Tempotron rule	91.3
Two layer network [36]	Spike-based	Unsupervised	STDP	93.5
Spiking RBM [49]	Rate-based	Supervised	Contrastive divergence	94.1
Two layer network [38]	Spike-based	Unsupervised	STDP	95.0
Convolutional SNN [50]	Rate-based	Supervised	Back-propagation	**99.1**
Proposed SDNN	Spike-based	Unsupervised	STDP	98.4

neuronal maps in the first convolutional layer. The green and red colors correspond to ON- and OFF-center DoG filters. Interestingly, this layer converged to Gabor-like edge detectors with different orientations, phase and polarity. These edge features are combined in the next layer and provide easily separable digit representation.

Recognition performance of the proposed method and some recent SNNs on the MNIST dataset are provided in Table 2. As seen, the proposed SDNN outperforms unsupervised SNNs by reaching 98.4% recognition accuracy. Besides, the accuracy of the proposed SDNN is close to the 99.1% accuracy of the totally supervised rate-based SDNN [50] which is indeed the converted version of a traditional DCNN trained by back-propagation.

The important advantage of the proposed SDNN is the use of much fewer spikes. Our SDNN uses only about 600 spikes for each MNIST images in total for all the layers, while the supervised rate-based SDNN uses thousands of spikes per layer [50]. Also, because of using rate-based neural coding in such networks, they need to process images for hundreds of time steps, while our network process the MNIST images in 30 time steps only. As stated in Section 2, the proposed SDNN uses a temporal code which encodes the information of the input image in the spike times, and each neuron in all layers, is allowed to fire at most once. This temporal code associated with unsupervised STDP rule leads to a fast, accurate, and efficient processing.

4 Discussion

Recent supervised DCNNs have reached high accuracies on the most challenging object recognition datasets such as Imagenet. Architecture of theses networks are largely inspired by the deep hierarchical processing in the visual cortex. For instance, DCNNs use retinotopically arranged neurons with restricted receptive fields, and the receptive field size and feature complexity of the neurons gradually increase through the layers. However, the learning and neural processing mechanisms applied in DCNNs are inconsistent with the visual cortex, where neurons communicate using spikes and learn the input spike patterns in a mainly unsupervised manner. Employing such mechanisms in DCNNs can improve their energy consumption and decrease their need for an expensive supervised learning with millions of labeled images.

A popular approach in previous researches is to convert pre-trained supervised DCNNs into equivalent spiking network. To simulate the floating-point calculations in DCNNs, they have to use the firing rate as the neural code, which in result increases the number of required spikes and the processing time. For instance, the converted version of a simple two-layer DCNN for the MNIST dataset with images of 28×28 pixels requires thousands of spikes and hundreds of time steps per image. On the other hand, there are some SDNNs [31, 32, 33] which are originally spiking network but employ firing rate coding and biologically implausible learning rules such as autoencoders and back-propagation.

Here, we proposed a STDP-based SDNN with a spike-time coding. Each neuron was allowed to fire at most once, where its spike-time indicates the

significance of its visual input. Therefore, neurons that fire earlier are carrying more salient visual information, and hence, they were allowed to do the STDP and learn the input patterns. As the learning progresses, each layer converged to a set of diverse but informative features, and the feature complexity gradually increases through the layers from simple edge features to object prototypes. The proposed SDNN was evaluated on several image datasets and reached high recognition accuracies. This shows how the proposed temporal coding and learning mechanism (STDP and learning competition) lead to discriminative object representations.

The proposed SDNN has several advantages to its counterparts. First, our proposed SDNN is the first spiking neural network with more than one learnable layer which can process large-scale natural object images. Second, due to the use of an efficient temporal coding, which encodes the visual information in the time of the first spikes, it can process the input images with a low number of spikes and in a few processing time steps. Third, the proposed SDNN exploits the bio-inspired and totally unsupervised STDP learning rule which can learn the diagnostic object features and neglect the irrelevant backgrounds.

Our SDNN could be efficiently implemented in parallel hardware (e.g., FPGA [51]) using address event representation (AER) [52] protocol. With AER, spike events are represented by the addresses of sending and receiving neurons, and time is represented by the asynchronous occurrence of spike events. Since these hardware are much faster than biological hardware, simulations could run several order of magnitude faster than real time [53]. The primate visual system extracts the rough content of an image in about 100 ms [3, 4, 54, 5]. We thus speculate that some dedicated hardware will be able to do the same in the order of a millisecond or less.

Also, the proposed SDNN can be modified to use spiking retinal models [55, 56] as the input layer. These models mimic the spatiotemporal filtering of the retinal ganglion cells with center/surround receptive fields. Alternatively, we could use neuromorphic asynchronous event-based cameras such as dynamic vision sensor (DVS), which generate output events when they capture transients in the scene [57]. Finally, due to the DoG filtering in the input layer of the proposed SDNN, some visual information such as texture and color are lost. Hence, future studies should focus on encoding these additional pieces of information in the input layer.

Biological evidence indicate that in addition to the unsupervised learning mechanisms (e.g., STDP), there are also dopamine-based reinforcement learning strategies in the brain [58]. Besides, although it is still unclear how supervised learning is implemented in biological neural networks, it seems that for some tasks (e.g., motor control and sensory inputs prediction) the brain must constantly learn temporal dynamics based on error feedback [59]. Employing such reinforcement and supervised learning strategies could improve the proposed SDNN in different aspects which are inevitable with unsupervised learning methods. Particularly, they can help to reduce the number of required features and to extract optimized task-dependent features.

Supplementary Information

Video S1. The learning progress and neural activity over the Caltech face/motorbike task. Here we presented the face and motorbike training examples, propagated the corresponding spike waves, and applied the STDP rule. The input image is presented at the top-left corner of the screen. The output spikes of the input layer (i.e., DoG layer) at each time step is presented in the top-middle panel, and the accumulation of theses spikes is shown in the top-right panel. For each of the subsequent convolutional layers, the preferred features, the output spikes at each time step, and the accumulation of the output spikes are presented in the corresponding panels. Note that 4, 8, and 2 features from the first, second and third convolutional layers are selected and shown, respectively. As mentioned, the learning occurs layer by layer, thus, the label of the layer which is currently doing the learning is specified by the red color. As seen, the first layer learns to detect edges, the second layer learns intermediate features, and finally

the third layer learns face and motorbike prototype features.
The video is available at `https://youtu.be/u32Xnz2hDkE`

Acknowledgment

This research received funding from the European Research Council under the European Unions 7th Framework Program (FP/2007-2013) / ERC Grant Agreement n.323711 (M4 project).

References

[1] J. J. DiCarlo, D. Zoccolan, N. C. Rust, How does the brain solve visual object recognition?, Neuron 73 (3) (2012) 415–434.

[2] J. J. DiCarlo, D. D. Cox, Untangling invariant object recognition, Trends in Cognitive Sciences 11 (8) (2007) 333–341.

[3] S. Thorpe, D. Fize, C. Marlot, et al., Speed of processing in the human visual system, Nature 381 (6582) (1996) 520–522.

[4] C. P. Hung, G. Kreiman, T. Poggio, J. J. DiCarlo, Fast readout of object identity from macaque inferior temporal cortex, Science 310 (5749) (2005) 863–866.

[5] H. Liu, Y. Agam, J. R. Madsen, G. Kreiman, Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex, Neuron 62 (2) (2009) 281–290.

[6] K. Fukushima, Neocognitron : a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position., Biological Cybernetics 36 (4) (1980) 193–202.

[7] Y. LeCun, Y. Bengio, Convolutional networks for images, speech, and time series, in: The Handbook of Brain Theory and Neural Networks, Cambridge, MA: MIT Press, 1998, pp. 255–258.

[8] T. Serre, L. Wolf, S. Bileschi, M. Riesenhuber, T. Poggio, Robust object recognition with cortex-like mechanisms, IEEE Transactions on Pattern Analysis Machine Intelligence 29 (3) (2007) 411–426.

[9] T. Masquelier, S. J. Thorpe, Unsupervised learning of visual features through spike timing dependent plasticity, PLoS Computational Biology 3 (2) (2007) e31.

[10] H. Lee, R. Grosse, R. Ranganath, A. Y. Ng, Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations, ACM Press, New York, New York, USA, 2009, pp. 1–8.

[11] N. Pinto, Y. Barhomi, D. D. Cox, J. J. DiCarlo, Comparing state-of-the-art visual features on invariant object recognition tasks, in: IEEE workshop on Applications of Computer Vision (WACV), Kona, Hawaii, USA, 2011, pp. 463–470.

[12] M. Ghodrati, A. Farzmahdi, K. Rajaei, R. Ebrahimpour, S.-M. Khaligh-Razavi, Feedforward object-vision models only tolerate small image variations compared to human, Frontiers in Computational Neuroscience 8 (74) (2014) 1–17.

[13] A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep convolutional neural networks., in: Neural Information Processing Systems (NIPS), Lake Tahoe, Nevada, USA, 2012, pp. 1–9.

[14] M. D. Zeiler, R. Fergus, Visualizing and understanding convolutional networks, in: European Conference on Computer Vision (ECCV), Zurich, Switzerland, 2014, pp. 818–833.

[15] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv:1409.1556.

[16] S. R. Kheradpisheh, M. Ghodrati, M. Ganjtabesh, T. Masquelier, Deep networks resemble human feedforward vision in invariant object recognition, Scientific Reports 6 (2016) 32672.

[17] S. R. Kheradpisheh, M. Ghodrati, M. Ganjtabesh, T. Masquelier, Humans and deep networks largely agree on which kinds of variation make object recognition harder, Frontiers in Computational Neuroscience 10 (2016) 92.

[18] C. F. Cadieu, H. Hong, D. L. Yamins, N. Pinto, D. Ardila, E. A. Solomon, N. J. Majaj, J. J. DiCarlo, Deep neural networks rival the representation of primate it cortex for core visual object recognition, PLoS Computational Biology 10 (12) (2014) e1003963.

[19] S.-M. Khaligh-Razavi, N. Kriegeskorte, Deep supervised, but not unsupervised, models may explain it cortical representation, PLoS Computational Biology 10 (11) (2014) e1003915.

[20] R. M. Cichy, A. Khosla, D. Pantazis, A. Torralba, A. Oliva, Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence, Scientific Reports 6.

[21] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (7553) (2015) 436–444.

[22] S. Shoham, D. H. OConnor, R. Segev, How silent is the brain: is there a dark matter problem in neuroscience?, Journal of Comparative Physiology A 192 (8) (2006) 777–784.

[23] W. Maass, Computing with spikes, Special Issue on Foundations of Information Processing of TELEMATIK 8 (1) (2002) 32–36.

[24] E. T. Rolls, G. Deco, Computational neuroscience of vision, Oxford university press, Oxford, UK, 2002.

[25] C. D. Meliza, Y. Dan, Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking, Neuron 49 (2) (2006) 183–189.

[26] S. Huang, C. Rozas, M. Treviño, J. Contreras, S. Yang, L. Song, T. Yoshioka, H.-K. Lee, A. Kirkwood, Associative hebbian synaptic plasticity in primate visual cortex, The Journal of Neuroscience 34 (22) (2014) 7575–7579.

[27] D. B. McMahon, D. A. Leopold, Stimulus timing-dependent plasticity in high-level vision, Current Biology 22 (4) (2012) 332–337.

[28] Y. Cao, Y. Chen, D. Khosla, Spiking deep convolutional neural networks for energy-efficient object recognition, International Journal of Computer Vision 113 (1) (2015) 54–66.

[29] E. Hunsberger, C. Eliasmith, Spiking deep networks with lif neurons, arXiv:1510.08829.

[30] P. U. Diehl, G. Zarrella, A. Cassidy, B. U. Pedroni, E. Neftci, Conversion of artificial recurrent neural networks to spiking neural networks for low-power neuromorphic hardware, in: IEEE International Conference on Rebooting Computing, San Diego, California, USA, 2016, pp. 1–8.

[31] P. Panda, K. Roy, Unsupervised regenerative learning of hierarchical features in spiking deep networks for object recognition, in: IEEE International Joint Conference on Neural Networks (IJCNN), Vancouver, Canada, 2016, pp. 1–8.

[32] K. S. Burbank, Mirrored stdp implements autoencoder learning in a network of spiking neurons, PLoS Computational Biology 11 (12) (2015) e1004566.

[33] Y. Bengio, D.-H. Lee, J. Bornschein, Z. Lin, Towards biologically plausible deep learning, arXiv:1502.04156.

[34] J. M. Brader, W. Senn, S. Fusi, Learning real-world stimuli in a neural network with spike-driven synaptic dynamics, Neural computation 19 (11) (2007) 2881–2912.

[35] S. Habenschuss, J. Bill, B. Nessler, Homeostatic plasticity in bayesian spiking networks as expectation maximization with posterior constraints, in: Advances in Neural Information Processing Systems, Lake Tahoe, Nevada, USA, 2012, pp. 773–781.

[36] D. Querlioz, O. Bichler, P. Dollfus, C. Gamrat, Immunity to device variations in a spiking neural network with memristive nanodevices, IEEE Transactions on Nanotechnology 12 (3) (2013) 288–295.

[37] B. Zhao, R. Ding, S. Chen, B. Linares-Barranco, H. Tang, Feedforward categorization on aer motion events using cortex-like features in a spiking neural network, IEEE Transactions on Neural Networks and Learning Systems 26 (9) (2015) 1963–1978.

[38] P. U. Diehl, M. Cook, Unsupervised learning of digit recognition using spike-timing-dependent plasticity, Frontiers in computational neuroscience 9 (2015) 99.

[39] M. Beyeler, N. D. Dutt, J. L. Krichmar, Categorization and decision-making in a neurobiologically plausible spiking network using a stdp-like learning rule, Neural Networks 48 (2013) 109–124.

[40] S. R. Kheradpisheh, M. Ganjtabesh, T. Masquelier, Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition, Neurocomputing 205 (2016) 382–392.

[41] S. Thorpe, A. Delorme, R. Van Rullen, Spike-based strategies for rapid processing, Neural Networks 14 (6) (2001) 715–725.

[42] R. Van Rullen, S. J. Thorpe, Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex, Neural Computation 13 (6) (2001) 1255–1283.

[43] G. Portelli, J. M. Barrett, G. Hilgen, T. Masquelier, A. Maccione, S. Di Marco, L. Berdondini, P. Kornprobst, E. Sernagor, Rank order coding: a retinal information decoding strategy revealed by large-scale multielectrode array retinal recordings, Eneuro 3 (3) (2016) ENEURO–0134.

[44] A. Delorme, L. Perrinet, S. J. Thorpe, Networks of integrate-and-fire neurons using rank order coding b: Spike timing dependent plasticity and emergence of orientation selectivity, Neurocomputing 38 (2001) 539–545.

[45] G. A. Rousselet, S. J. Thorpe, M. Fabre-Thorpe, Taking the max from neuronal responses, Trends in Cognitive Sciences 7 (3) (2003) 99–102.

[46] F. Zou, Y. Wang, Y. Yang, K. Zhou, Y. Chen, J. Song, Supervised feature learning via l 2-norm regularized logistic regression for 3d object recognition, Neurocomputing 151 (2015) 603–611.

[47] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324.

[48] S. Hussain, S.-C. Liu, A. Basu, Improved margin multiclass classification using dendritic neurons with morphological learning, in: IEEE International Symposium on Circuits and Systems (ISCAS), Melbourne, VIC, Australia, 2014, pp. 2640–2643.

[49] P. O'Connor, D. Neil, S.-C. Liu, T. Delbruck, M. Pfeiffer, Real-time classification and sensor fusion with a spiking deep belief network, Frontiers in Neuroscience 7 (2013) 178.

[50] P. U. Diehl, D. Neil, J. Binas, M. Cook, S.-C. Liu, M. Pfeiffer, Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing, in: IEEE International Joint Conference on Neural Networks (IJCNN), Killarney, Ireland, 2015, pp. 1–8.

[51] A. Yousefzadeh, T. Serrano-Gotarredona, B. Linares-Barranco, Fast Pipeline 128??128 pixel spiking convolution core for event-driven vision processing in FPGAs, in: Proceedings of 1st International Conference on Event-Based Control, Communication and Signal Processing (EBCCSP), 2015.

[52] M. Sivilotti, Wiring considerations in analog VLSI systems with application to field-programmable networks, Ph.D. thesis, Comput. Sci. Div., California Inst. Technol., Pasadena, CA (1991).

[53] T. Serrano-Gotarredona, T. Masquelier, T. Prodromakis, G. Indiveri, B. Linares-Barranco, STDP and STDP variations with memristors for spiking neuromorphic learning systems., Frontiers in Neuroscience 7 (2013) 2.

[54] H. Kirchner, S. J. Thorpe, Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited, Vision Research 46 (11) (2006) 1762–1776.

[55] A. Wohrer, P. Kornprobst, Virtual Retina: a biological retina model and simulator, with contrast gain control. 26 (2) (2009) 219–49.

[56] P. Martínez-Cañada, C. Morillas, B. Pino, E. Ros, F. Pelayo, A Computational Framework for Realistic Retina Modeling, International Journal of Neural Systems 26 (7) (2016) 1650030.

[57] P. Lichtsteiner, C. Posch, T. Delbruck, An 128x128 120dB 15us-latency temporal contrast vision sensor, IEEE J. Solid State Circuits 43 (2) (2007) 566–576.

[58] M. Pignatelli, A. Bonci, Role of dopamine neurons in reward and aversion: a synaptic plasticity perspective, Neuron 86 (5) (2015) 1145–1157.

[59] K. Doya, Complementary roles of basal ganglia and cerebellum in learning and motor control, Current Opinion in Neurobiology 10 (6) (2000) 732–739.